

Efficient Smoothing of 3-D Unstructured Grids

Markus Berndt and J. David Moulton, T-7;
and G. A. Hansen, South Dakota School of
Mines and Technology

The smoothing of unstructured grids in large-scale three-dimensional simulations is a critical component of many projects related to the Advanced Simulation and Computing Project. Few of the existing mesh smoothing techniques are capable of handling large unstructured meshes in complex geometries. For example, most smoothers generate unacceptable or even invalid grids in the neighborhood of extremely convex or concave boundaries. In contrast, a smoothing technique based on harmonic coordinates was developed in [1] that is robust with respect to these geometric complexities. This technique defines a system of three quasi-linear diffusion equations, one for each coordinate direction. The only coupling between components of this system is through the elements of the metric tensor, which defines the solution-dependent diffusion tensor. In turn, grid smoothing is driven by a target metric tensor that is based on a coarse-graining (local averaging) of the current mesh. We use a standard vertex-based Finite Element Method (FEM) to discretize the variational formulation of this quasi-linear system of equations.

We have developed a Jacobian-Free Newton-Krylov (JFNK) algorithm [2] with a multilevel preconditioner to solve this discrete nonlinear system of equations. The JFNK approach recognizes that the underlying Krylov iteration (e.g., GMRES) does not require the Jacobian matrix itself, but simply the action of the Jacobian on a vector. Since this action may be approximated through two nonlinear function evaluations, it is not necessary to compute the Jacobian. However,

without effective preconditioning the performance of this GMRES iteration will be unacceptable for large meshes. In particular, to achieve optimal algorithmic scaling we must leverage multilevel solution techniques in the preconditioner, and these techniques require a matrix.

Thus, the fundamental challenge is to develop a preconditioner that is inexpensive to form, effective in its approximation of the Jacobian, and amenable to multilevel solvers. To motivate the preconditioners that we investigated, we note that the Jacobian may be written as the sum of a symmetric block diagonal operator \mathcal{D} (the Picard linearization) and a nonsymmetric operator \mathcal{N} with no derivative terms.

$$\mathcal{J} = \mathcal{D} + \mathcal{N}$$

The properties of \mathcal{J} differ significantly from the underlying assumptions and supporting theory found in the most widely used class of multilevel iterative solvers, namely, algebraic multigrid (AMG). Nevertheless, AMG can be used to solve these systems, and although its performance is somewhat less predictable and its algorithmic scaling is suboptimal, it is significantly more efficient than a single level preconditioner. Hence, our first JFNK algorithm forms the full Jacobian only once for the initial mesh, and then reuses it throughout the nonlinear iteration. This *frozen Jacobian* preconditioner is treated with AMG.

However, to deliver the most efficient solver possible we must eliminate the evaluation of the Jacobian entirely, and replace it with an accurate and inexpensive approximation. The symmetric block diagonal Picard linearization, \mathcal{D} , is ideal from this perspective. There are three identical blocks in \mathcal{D} that correspond to a scalar diffusion problem that is effectively treated with AMG.

We illustrate the excellent performance of our JFNK solvers [3] through a scaling study that employs a sequence of refined meshes based on the object in Fig. 1. The poor scaling and high computational cost of the original solver is shown in Fig. 2. In contrast, our JFNK algorithm with GMRES preconditioned by AMG-V-cycles converged to a relative residual of 0.01 on the *frozen Jacobian* decreases the computation time by approximately a factor of five for the largest mesh in this study, although it scales poorly. Finally, using the Picard linearization, \mathcal{D} , as the preconditioner eliminates the need for any Jacobian evaluations. Our JFNK implementation with GMRES preconditioned by AMG-V-cycles converged to a relative residual of 0.01 on \mathcal{D} achieves an impressive 20x speedup over the original algorithm for the largest mesh in this study and the best algorithmic scaling.

For more information contact J. David Moulton at moulton@lanl.gov.

- [1] G. Hansen, et al., *J. Comp. Phys.* **202** (1), 281–297 (2005).
- [2] D.A. Knoll and D.E. Keyes, *J. Comp. Phys.* **193** (2) 357–397 (2004).
- [3] M. Berndt, et al., “A Scalable Non-linear Solver for the Laplace-Beltrami Mesh Generation Equations,” Los Alamos National Laboratory report LA-UR-05-6189 (2005).

Funding Acknowledgements

NNSA’s Advanced Simulation and Computing (ASC), Advanced Applications Program.




Fig. 1. An object with a re-entrant crease (top) is meshed initially with transfinite interpolation. Many elliptic smoothing techniques (e.g., Winslow) fold the mesh along the re-entrant crease. The target metric approach provides a well-balanced distribution of cell volumes and angles near this line (bottom). Similar improvements are observed for common awkward internal structures, such as triple-points.




Fig. 2. A scaling study based on three successively finer meshes of the object in Fig. 1, illustrates the gain in performance resulting from the new Jacobian-Free Newton-Krylov (JFNK) solver with advanced multilevel preconditioners. We observe that for the largest mesh, JFNK preconditioned with AMG V-cycles on the Picard linearization (converged to a relative residual of 0.01) is nearly 20 times faster than the original NK solver that used a block ILU preconditioner.